

On the zone of a circle in an arrangement of lines*

Gabriel Nivasch[†]

Abstract

Let \mathcal{L} be a set of n lines in the plane, and let C be a convex curve in the plane, like a circle or a parabola. The *zone* of C in \mathcal{L} , denoted $\mathcal{Z}(C, \mathcal{L})$, is defined as the set of all cells in the arrangement $\mathcal{A}(\mathcal{L})$ that are intersected by C . Edelsbrunner et al. (1992) showed that the complexity (total number of edges or vertices) of $\mathcal{Z}(C, \mathcal{L})$ is at most $O(n\alpha(n))$, where α is the inverse Ackermann function. They did this by translating the sequence of edges of $\mathcal{Z}(C, \mathcal{L})$ into a sequence S that avoids the subsequence *ababa*. Whether the worst-case complexity of $\mathcal{Z}(C, \mathcal{L})$ is only linear is a longstanding open problem.

Since the relaxation of the problem to pseudolines does have a $\Theta(n\alpha(n))$ bound, any proof of $O(n)$ for the case of straight lines must necessarily use geometric arguments.

In this paper we present some such geometric arguments. We show that, if C is a circle, then certain configurations of straight-line segments with endpoints on C are impossible. In particular, we show that there exists a Hart–Sharir sequence that cannot appear as a subsequence of S .

The Hart–Sharir sequences are essentially the only known way to construct *ababa*-free sequences of superlinear length. Hence, if it could be shown that every family of *ababa*-free sequences of superlinear-length eventually contains all Hart–Sharir sequences, it would follow that the complexity of $\mathcal{Z}(C, \mathcal{L})$ is $O(n)$ whenever C is a circle.

1 Introduction

Let \mathcal{L} be a set of n lines in the plane. The *arrangement* of \mathcal{L} , denoted $\mathcal{A}(\mathcal{L})$, is the partition of the plane into vertices, edges, and cells induced by \mathcal{L} . Let C be another object in the plane. The *zone* of C in \mathcal{L} , denoted $\mathcal{Z}(C, \mathcal{L})$, is defined as the set of all cells in $\mathcal{A}(\mathcal{L})$ that are intersected by C . The *complexity* of $\mathcal{Z}(C, \mathcal{L})$ is defined as the total number of edges, or vertices, in it.

The celebrated *zone theorem* states that, if C is another line, then $\mathcal{Z}(C, \mathcal{L})$ has complexity $O(n)$ (Chazelle et al. [4]; see also Edelsbrunner et al. [6], Matoušek [15]).

If C is a convex curve, like a circle or a parabola, then $\mathcal{Z}(C, \mathcal{L})$ is known to have complexity $O(n\alpha(n))$, where α is the very-slow-growing inverse Ackermann function (Edelsbrunner et al. [6]; see also Bern et al. [3], Sharir and Agarwal [24]). More specifically, the *outer zone* of $\mathcal{Z}(C, \mathcal{L})$ (the part that lies outside the convex hull of C) is known to have complexity $O(n)$, whereas the complexity of the *inner zone* is only known to be $O(n\alpha(n))$. Whether the complexity of the inner zone is linear as well is a longstanding open problem [3, 24].

*An extended abstract of this paper appeared in *EuroComb 2015 (Electronic Notes in Discrete Mathematics 49:221–231, 2015)*.

[†]gabrieln@ariel.ac.il. Department of Computer Science, Ariel University, Ariel, Israel.

The gap between the upper and the lower bound is completely negligible for all practical purposes, but the question is interesting from a purely mathematical point of view.

In this paper we make progress towards proving that the inner zone of a circle in an arrangement of lines has linear complexity. Since we find it easier to work with a parabola than with a circle, throughout this paper we will take C to be the parabola $y = x^2$. The two problems are equivalent by a projective transformation, as we will explain.

1.1 Davenport–Schinzel sequences and their generalizations

Let S be a finite sequence of symbols, and let $s \geq 1$ be a parameter. Then S is called a *Davenport–Schinzel sequence of order s* if every two adjacent symbols in S are distinct, and if S does not contain any alternation $a \cdots b \cdots a \cdots b \cdots$ of length $s + 2$ for two distinct symbols $a \neq b$. Hence, for $s = 1$ the “forbidden pattern” is aba , for $s = 2$ it is $abab$, for $s = 3$ it is $ababa$, and so on.

The maximum length of a Davenport–Schinzel sequence of order s that contains only n distinct symbols is denoted $\lambda_s(n)$. For $s \leq 2$ we have $\lambda_1(n) = n$ and $\lambda_2(n) = 2n - 1$. However, for fixed $s \geq 3$, $\lambda_s(n)$ is slightly superlinear in n .

DS sequences of order 3 The case $s = 3$ is the one most relevant to us. Hart and Sharir [9] (see also [16, 24]) constructed a family of sequences that achieve the lower bound¹ $\lambda_3(n) \geq n\alpha(n) - O(n)$; and they also proved the asymptotically matching upper bound $\lambda_3(n) \leq O(n\alpha(n))$. Klazar [12] subsequently improved the upper bound to $\lambda_3(n) \leq 2n\alpha(n) + O(n\sqrt{\alpha(n)})$ (recently, Pettie [21] improved the lower-order term to $O(n)$).

Nivasch [16] showed that $\lambda_3(n) \geq 2n\alpha(n) - O(n)$. Hence, $\lambda_3(n) = 2n\alpha(n) \pm O(n)$. Nivasch’s construction is an *extension* of the Hart–Sharir construction, in the sense that Nivasch’s sequences contain the Hart–Sharir sequences as subsequences.² Geneson [8] made a nice cosmetic improvement to Nivasch’s construction.

DS sequences of higher orders For $s = 4$ we have $\lambda_4(n) = \Theta(n \cdot 2^{\alpha(n)})$, and in general, $\lambda_s(n) = n \cdot 2^{\Theta(\text{poly}(\alpha(n)))}$ for fixed $s \geq 4$, where the polynomial in the exponent is of degree roughly $s/2$. See Sharir and Agarwal [24], and subsequent improvements by Nivasch [16] and Pettie [21].

Generalized DS sequences A *generalized Davenport–Schinzel sequence* is one where the forbidden pattern is not restricted to be $abab \cdots$, but it can be any fixed subsequence u . In order for the problem to be nontrivial we must require S to be k -sparse—meaning, every k adjacent symbols in S must be pairwise distinct—where $k = \|u\|$ is the number of distinct symbols in u . For example, if we take $u = abcacbc$, then S must not contain any subsequence of the form $a \cdots b \cdots c \cdots a \cdots c \cdots c \cdots b \cdots c$ for $|\{a, b, c\}| = 3$, and every three adjacent symbols in S must be pairwise distinct.

We denote by $\text{Ex}(u, n)$ the maximum length of a k -sparse, u -avoiding sequence S on n distinct symbols, where $k = \|u\|$. For every fixed forbidden pattern u , $\text{Ex}(u, n)$ is at most

¹The bound claimed in [24] is $\lambda_3(n) \geq (1/2)n\alpha(n) - O(n)$, because a factor of 2 is lost in interpolation; this problem is fixed in [16].

²This can be shown with an argument similar to that of Lemma 5.7 below, which is beyond the scope of this paper.

slightly superlinear in n : $\text{Ex}(u, n) = O(n \cdot 2^{\text{poly}(\alpha(n))})$, where the polynomial in the exponent depends on u (Klazar [10], Nivasch [16], Pettie [22]).

Similarly, if $U = \{u_1, u_2, \dots, u_j\}$ is a set of patterns, then $\text{Ex}(U, n)$ denotes the maximum length of a sequence that avoids all the patterns in U , is k -sparse for $k = \min\{\|u\| : u \in U\}$, and contains only n distinct symbols.

Some relevant results on generalized DS sequences Let us mention some results on generalized DS sequences that are relevant to us:

- $\text{Ex}(\{ababa, abcacbc\}, n) = \Theta(n\alpha(n))$ (Pettie [19]). Indeed, the *ababa*-free sequences of Hart and Sharir [9] avoid *abcacbc* as well.³ See Section 5 below.
- $\text{Ex}(abcacbc, n) = \Theta(n\alpha(n))$ (Pettie [20]). The lower bound is achieved by a modification of the Hart–Sharir construction, which does not avoid *ababa* anymore.
- It is unknown whether $\text{Ex}(\{ababa, abcacbc\}, n)$ or $\text{Ex}(\{ababa, abcacbc, (abcacbc)^R\}, n)$ are superlinear in n (where u^R denotes the reversal of u). We conjecture that they are both $O(n)$.

Applications of generalized DS sequences Generalized Davenport–Schinzel sequences have found a few applications. Cibulka and Kynčl [5] used them to bound the size of sets of permutations with bounded VC-dimension. Valtr [26], Fox et al. [7], and Suk and Walczak [25] have used Generalized DS sequences to bound the number of edges in graphs with no k pairwise crossing edges; the papers [26, 7] use the “ N -shaped” forbidden pattern $a_1 \cdots a_\ell \cdots a_1 \cdots a_\ell$, and the papers [7, 25] use the forbidden pattern $(a_1 \cdots a_\ell)^m$.

Pettie considered $\text{Ex}(\{abababa, abaabba\}, n)$ for analyzing the deque conjecture for splay trees [17], and $\text{Ex}(\{ababab, abbaabba\}, n)$ for analyzing the union of fat triangles in the plane [18].

1.2 Transcribing the zone into a Davenport–Schinzel sequence

Let \mathcal{L} be a set of n lines in the plane, and let C be a convex curve in the plane. We can assume without loss of generality that C is either closed (like a circle) or unbounded in both directions (like a parabola), by prolonging C if necessary. Thus, C divides the plane into two regions, one of which equals the convex hull of C .

Here we recall the argument of Edelsbrunner et al. [6] showing that the complexity of the part of $\mathcal{Z}(C, \mathcal{L})$ that lies inside the convex hull of C is $O(n\alpha(n))$.

If C is unbounded in both directions then assume without loss of generality that it is x -monotone and it is the graph of a convex function, by rotating the whole picture if necessary.

Also assume general position for simplicity: No line of \mathcal{L} is vertical, no two lines are parallel, no three lines are concurrent, no line is tangent to C , and no two lines intersect C at the same point. (Perturbing \mathcal{L} into general position can only increase the complexity of $\mathcal{Z}(C, \mathcal{L})$.) We can also assume that every line of \mathcal{L} intersects C , since otherwise the line would not contribute to the complexity of the inner zone of C .

Let \mathcal{L}' be the set of n segments obtained by intersecting each line of \mathcal{L} with the convex hull of C . (If C is unbounded then some elements of \mathcal{L}' may actually be rays.)

³Spaces are just for clarity. The Hart–Sharir construction also avoids other patterns, such as *abcbdadbc* (Klazar [11]; see Pettie [19]).

Let G be the *intersection graph* of \mathcal{L}' , i.e. the graph having \mathcal{L}' as vertex set, and having an edge connecting two elements of \mathcal{L}' if and only if they intersect. We can assume without loss of generality that G is connected: If G has several connected components, then we can separately bound the complexity produced by each one and add them up; this works because our desired bound is at least linear in n . Since G is connected, all the bounded cells of the inner zone are simple (i.e. they touch C in a single interval); and if C is unbounded then there are at most two upward-unbounded cells (bounded by the two infinite extremes of C).

If C is closed then let c_0 be the topmost point of C ; we will ignore the cell that contains c_0 , since it has at most linear complexity (as any single cell does). If C is unbounded then we will similarly ignore the up-to-two unbounded cells.

To bound the complexity of the remaining cells, we will traverse their boundary and transcribe it into a sequence in a certain way.

Every segment of \mathcal{L}' has two sides, one of which will be called *positive* and the other one *negative*, as follows: If C is closed, then the positive side is the one facing the point c_0 and the negative side is the other one; if C is unbounded, then the positive side is the upper one and the negative side is the lower one.

If C is closed, then let c_1 be the first endpoint of \mathcal{L}' counterclockwise from c_0 along C , and let c_2 be the last endpoint. If C is unbounded, then c_1 is defined as the leftmost endpoint of \mathcal{L}' , and c_2 as the rightmost endpoint.

We traverse the boundary of the inner zone of C by starting at c_1 , and walking around the boundary of the cells, as if the segments were walls which we touch with the left hand at all times, until we reach c_2 . See Figure 1. We transcribe this tour into a sequence containing $3n$ distinct symbols as follows:

Each segment $a \in \mathcal{L}'$ is partitioned by the other segments into smaller pieces. We take two directed copies of each such piece. We call each such copy a *sub-segment*. The sub-segments are directed counterclockwise around a ; i.e. those above a are directed leftwards, and those below a are directed rightwards. Hence, our tour visits some of these sub-segments, in the directions we have given them, in a certain order.

For each segment a , the sub-segments of a that are visited, are visited in counterclockwise order around a . We first visit some sub-segments on the positive side of a , then we visit some sub-segments on the negative side of a , and then we again visit some sub-segments on the positive side of a .

Sub-segments of the first type are transcribed as a' ; sub-segments of the second type are transcribed as a , and sub-segments of the third type are transcribed as a'' . See again Figure 1. Let S' be the sequence resulting from the tour.

For each segment a , label its endpoints L_a and R_a , such that L_a is visited before R_a .⁴

Let a, b be two intersecting segments, such that L_a is visited before L_b . Then the restriction of S' to $\{a', a, a'', b', b, b''\}$ is of the form

$$(a')^* a^* (b')^* b^* a^* (a'')^* b^* (b'')^* (a'')^* \quad \text{or} \quad (b')^* (a')^* a^* (b')^* b^* a^* (a'')^* b^* (b'')^*,$$

where $*$ denotes zero or more repetitions. See Figure 2.

Hence, the restriction of S' to first-type symbols contains no alternation $abab$, and it contains no adjacent repetitions either, as can be easily seen. Hence, it is an order-2 DS-sequence and so it has linear length. The same is true for the restriction of S' to third-type symbols.

⁴If C is unbounded then L_a is always the left endpoint of a .




Figure 1: Traversing the boundary of the inner zone of C .




Figure 2: Symbol alternations produced by two intersecting segments.

Thus, the important part of the sequence S' is its restriction to second-type symbols—those corresponding to the negative side of the segments. From now on we denote this subsequence S , and we call it the *lower inner-zone sequence of $\mathcal{Z}(C, \mathcal{L})$* .⁵ The sequence S contains no alternation *ababa*, and it contains no adjacent repetitions, as can be easily seen. Hence, S is an order-3 DS-sequence, and hence its length is at most $O(n\alpha(n))$.

1.3 Relation to lower envelopes

Lower envelopes are the original motivation for Davenport–Schinzel sequences. If $\mathcal{F} = \{f_1, \dots, f_n\}$ is a collection of n x -monotone curves in the plane (continuous functions $\mathbb{R} \rightarrow \mathbb{R}$), then the *lower envelope* of \mathcal{F} is their pointwise minimum (or the part that can be seen from the point $(0, -\infty)$), and the *lower-envelope sequence* is the sequence of functions that appear in the lower envelope, from left to right. If the f_i 's are partially defined functions (say, each one is defined only on an interval of \mathbb{R}), then the definition is the same, except that the symbol “ ∞ ” might also appear in the lower-envelope sequence.

In our case, if C is x -monotone, then the lower-envelope sequence of the set of segments \mathcal{L}' is a subsequence of S : It contains only those parts that can be seen from $-\infty$. We shall denote this sequence by $N = N(\mathcal{L}')$.

The Hart–Sharir sequences can be realized as lower-envelope sequences of segments in the plane (Wiernik and Sharir [27]; see also [15, 24]). However, it is unknown whether this is still possible if all the endpoints are required to lie on a circle/parabola (like our set \mathcal{L}'), or more generally on a convex curve. Sharir and Agarwal raise this question in [24, p. 112]. Proving a linear upper bound for the length of N might be easier than for the length of S .

It is also not known whether the longer sequences of Nivasch [16] can be realized as lower-envelope sequences of segments. It is not even known whether there *exists* an order-3 DS sequence that cannot be realized as a lower-envelope sequence of segments.

1.4 From circles to conic sections

Let us return to the zone problem. The special cases in which C is a circle, a parabola, or a hyperbola, are all equivalent, as can be shown by suitable projective transformations: Let $\pi_1 \subset \mathbb{R}^3$ be a plane that contains a set of lines \mathcal{L} and a circle C . Let K be a cone in \mathbb{R}^3 that intersects π_1 at C , and let π_2 be a plane that intersects K at a parabola C' . Then, the projection through the apex of K maps π_1 (with the exception of one line within π_1) into π_2 , mapping lines into lines, and mapping C (except for one point $p \in C$) into C' . We just have to choose π_2 so that no line of \mathcal{L} passes through p .

More concretely, the projective transformation $(x, y) \mapsto (\frac{x}{1-y}, \frac{1+y}{1-y})$ maps the unit circle $x^2 + y^2 = 1$ (except for the point $p = (0, 1)$) into the parabola $y = x^2$, mapping lines into lines.

The case of a hyperbola is handled similarly. First, note that all hyperbolas are equivalent under affine transformations. Hence, choose π_2 so that it intersects the cone K at a hyperbola C' , such that almost all of C is mapped to one branch of C' , and only a tiny portion of C , which does not intersect any line of \mathcal{L} , is mapped to the other branch of C' .

Even though the most natural formulation of the problem involves a circle, in this paper we will work with a parabola, since we find it easier to work with.

⁵Slight abuse of terminology. We will mainly deal with the case where C is unbounded; in this case the negative side of a segment is always its lower side.




Figure 3: Left: Realizing the sequence $u = abacdcbcd$ as a lower-envelope sequence of pseudosegments. Note that in this case the technique produces a supersequence of u . Right: Adding a convex curve and prolonging the pseudosegments into pseudolines.

1.5 The case of pseudolines and the need for geometric arguments

If we relax the problem and allow \mathcal{L} to consist of x -monotone *pseudolines* (x -monotone curves that pairwise intersect at most once and intersect C at most twice), then $\mathcal{Z}(C, \mathcal{L})$ can have complexity $\Theta(n\alpha(n))$. Indeed, in this setting *every* order-3 DS-sequence can appear as a subsequence of a lower-envelope sequence $N(\mathcal{L}')$. To see this, first note that every order-3 DS-sequence can appear as a subsequence of a lower-envelope sequence of x -monotone pseudosegments [24]; see Figure 3 for an example. Furthermore, in this construction, all segment endpoints are visible from $-\infty$. Hence, we can enclose the construction in a circle C , and prolong each pseudosegment ℓ into an x -monotone pseudoline by adding two very steeply decreasing rays on the two sides of ℓ .

Therefore, if, as we conjecture, the bound for the case of straight lines is only $O(n)$, then any proof must necessarily use geometric arguments, and not merely combinatorial ones.

1.6 Our results

In this paper we offer some evidence for the following conjecture, and make some progress towards proving it:

Conjecture 1.1. *If \mathcal{L} is a set of n lines and C is a circle, then the lower inner-zone sequence S of $\mathcal{Z}(C, \mathcal{L})$ has length $O(n)$, and hence $\mathcal{Z}(C, \mathcal{L})$ has at most linear complexity.*

Our technique consists of first finding segment configurations that are geometrically impossible, and then finding *ababa*-free sequences that *force* these configurations. We say that an *ababa*-free sequence u *forces* a segment configuration T , if, for every family of segments \mathcal{L}' (as in Section 1.2) whose lower inner-zone sequence contains u as a subsequence, \mathcal{L}' contains a subfamily combinatorially equivalent to T .

Thus, we first show in Section 3 that a certain, relatively simple configuration of eleven segments is impossible. Then we show that this configuration is forced by a pattern u of length 33. It follows that the lower inner-zone sequence S avoids u . This result, however, is useless for establishing Conjecture 1.1, since u contains both *abacacbc* and its reversal. Therefore, by the above-mentioned result of Pettie, the Hart–Sharir construction avoids both u and u^R (which is actually the same as u), and so $\text{Ex}(\{ababa, u, u^R\}, n) = \Theta(n\alpha(n))$.

Section 3 is just a warmup for Sections 4 and 5. In Section 4 we construct another impossible segment configuration X , this time with 173 segments. We could construct a

pattern u' that, as a consequence, cannot occur in S , but we abstain from doing so. Instead, we show directly in Section 5 that the Hart–Sharir sequences eventually force the configuration X .

Our results in Sections 4 and 5 were obtained as follows: Matoušek [15] and Sharir and Agarwal [24] describe a construction by P. Shor of segments in the plane whose lower-envelope sequences are the Hart–Sharir sequences. We tried to force the segment endpoints in the construction to lie on the parabola C , and we reached a contradiction. The impossible configuration X of Section 4 is the best way we found to isolate the contradiction. We elaborate more on this at the end of Section 5.

Next, in Section 6 we present some directions for further work on the problem: We formulate a conjecture regarding generalized DS sequences which, if true, would imply Conjecture 1.1. We also explain how our conjecture relates to previous research on generalized DS sequences.

Finally, in Section 7 we conclude by listing some related open problems.

2 Preliminaries

Throughout this and the following sections C will denote the parabola $y = x^2$, \mathcal{L}' will denote a set of n segments with endpoints on C , and S will denote the corresponding lower inner-zone sequence. As we said, we assume that no two segments have the same endpoint, and that the intersection graph of \mathcal{L}' is connected.

Recall that the left and right endpoints of a segment $a \in \mathcal{L}'$ are denoted L_a and R_a , respectively. Whenever we say that a sequence of endpoints appear in a certain order, we mean from left to right.

Let u be an *ababa*-free sequence in which, for simplicity, each symbol appears at least twice. Then, we define its *endpoint sequence* $E(u)$ by replacing, for each symbol a in u , its first occurrence by L_a and its last occurrence by R_a , and deleting all other occurrences of a . For example,

$$E(abacadbdbd) = L_a L_b L_c R_a L_d R_c R_b R_d.$$

It is clear that, if S is the lower inner-zone sequence of \mathcal{L}' , then the order of the endpoints of \mathcal{L}' is exactly $E(S)$. However, if u is a subsequence of S , then $E(u)$ is not necessarily a subsequence of $E(S)$; meaning, $E(S)$ does not necessarily respect the order of the symbols in $E(u)$. For example, if $u = abba$, so the order of the left endpoints in $E(u)$ is L_a, L_b , it is still possible for their order in $E(S)$ to be L_b, L_a .

Nevertheless, we now give a sufficient condition for guaranteeing that $E(u)$ is a subsequence of $E(S)$:

Definition 2.1. A symbol a in a sequence u is said to be *left-clamped* if u contains the subsequence $baba$, where b is the symbol immediately preceding the first a in u . Similarly, the symbol a is *right-clamped* if u contains the subsequence $abab$, where b is the symbol immediately following the last a in u .

Lemma 2.2. *If all the symbols in u , except for the very first symbol, are left-clamped, and all the symbols except for the very last symbol are right-clamped, and if S is an *ababa*-free supersequence of u , then $E(u)$ is a subsequence of $E(S)$.*




Figure 4: Left: Segments intersecting concavely. Right: Segments intersecting convexly.

Proof. Consider an adjacent pair of symbols Q_a, Q'_b in $E(u)$, where each of Q, Q' is either L or R . We claim that their order in $E(S)$ is also Q_a, Q'_b .

If Q is L and Q' is R , then trivially L_a also precedes R_b in $E(S)$.

If Q is R , then let c be the symbol immediately following the last a in u (note that c is not necessarily b). Since a is right-clamped, S cannot contain any a after the occurrence of b that gives rise to Q'_b ; otherwise, S would contain the forbidden pattern $acaca$.

Similarly, if Q' is L , let d be the symbol immediately preceding the first b in u . Since b is left-clamped, S cannot contain any b before the occurrence of a that gives rise to Q_a ; otherwise, S would contain $bdbdb$.

Hence, whether (Q, Q') equals (R, L) , (R, R) , or (L, L) , there is no way for Q_a, Q'_b to change order in $E(S)$. \square

Two segments a, b intersect if and only if their endpoints appear in the order $L_a L_b R_a R_b$ or $L_b L_a R_b R_a$.

If a_1, \dots, a_m are segments whose endpoints appear in the order $L_{a_1} \dots L_{a_m} R_{a_1} \dots R_{a_m}$, then they pairwise intersect. If the intersection points $a_m \cap a_{m-1}, \dots, a_3 \cap a_2, a_2 \cap a_1$ appear in this order from left to right, then we say that the segments *intersect concavely*. If the intersection points appear in the reverse order, then we say that the segments *intersect convexly*. See Figure 4.

If the segments a_1, a_2, \dots, a_n intersect concavely (or convexly), and $1 \leq i_1 < i_2 < \dots < i_k \leq n$ are increasing indices, then $a_{i_1}, a_{i_2}, \dots, a_{i_k}$ also intersect concavely (or convexly).

Observation 2.3. *If S contains the “N-shaped” subsequence $12 \dots m \dots 212 \dots m$, then the corresponding segments must have endpoints in the order $L_1 \dots L_m R_1 \dots R_m$, and must intersect concavely.*

Proof sketch. The general case follows from the case $m = 3$. \square

In Sections 3 and 4, we will specify some *segment configurations* by listing the order of their endpoints, and by specifying that some subsets of segments must intersect concavely. We will prove that some segment configurations are geometrically impossible.

2.1 Geometric properties of the parabola

We now present some simple geometric properties of the parabola C and straight-line segments. These properties lie at the heart of the distinction between the cases of straight lines and pseudolines, as we explained in the Introduction.

Observation 2.4. *Let $a, b \in \mathbb{R}$ be fixed. Then the affine transformation $m : \mathbb{R}^2 \rightarrow \mathbb{R}^2$ given by $m(x, y) = (ax + b, 2abx + a^2y + b^2)$ maps the parabola C to itself and keeps vertical lines*




Figure 5: Left: We have $p/q = r/s$. Right: Three segments intersecting concavely.

vertical. Therefore, we are free to horizontally translate and scale the set of x -coordinates of the segment endpoints \mathcal{L}' , without affecting the resulting lower inner-zone sequence S or the lower-envelope sequence N .⁶

Lemma 2.5. *Let a, b, c, d be four points on the parabola C , having increasing x -coordinates $a_x < b_x < c_x < d_x$. Let $z = ac \cap bd$. Define the horizontal distances $p = b_x - a_x$, $q = d_x - c_x$, $r = z_x - b_x$, $s = c_x - z_x$. Then $p/q = r/s$. See Figure 5, left.*

Proof. This can be shown directly by a slightly cumbersome algebraic calculation.

An alternative, more insightful proof uses elementary geometry and a limiting argument:

Let C be not a parabola but a unit circle. Let α be a very small angle, and let C_1 be the arc of C measuring angle α that is centered around the lowest point of C . Let a, b, c, d be four points on C_1 , in this order from left to right, and let $z = ac \cap bd$. Then, by the intersecting chords theorem, we have $ab/cd = bz/cz$. Since all the considered segments are almost horizontal, their length is almost equal to their x -projection. If we affinely stretch C_1 horizontally and vertically so its bounding box has width 1 and height 1, then it will almost match a parabola: At the limit as $\alpha \rightarrow 0$, the stretched arc pointwise converges to a parabolic segment. This affine transformation preserves the ratios between the horizontal projections, so the result follows. \square

Lemma 2.5 is actually part of a more general correspondence between circles and parabolas; see Yaglom [28].

Lemma 2.6. *Let a, b, c, d, e, f be six points on the parabola C , listed by increasing x -coordinate. Suppose the segments ad, be, cf intersect concavely. Define $\alpha_1 = b_x - a_x$, $\alpha_2 = c_x - b_x$, $\gamma = d_x - c_x$, $\beta_1 = e_x - d_x$, $\beta_2 = f_x - e_x$. Then:*

1. $\alpha_1/\beta_1 > \alpha_2/\gamma$ and $\beta_2/\alpha_2 > \beta_1/\gamma$;
2. $\alpha_1/\beta_1 > \alpha_2/\beta_2$;
3. $\beta_1 < \beta_2$ or $\alpha_2 < \gamma + \beta_1 + \beta_2$ (or both).

⁶This observation would be useful in *lower-bound constructions*, and hence, it is not used in the paper. We include it here just for the sake of completeness.




Figure 6: Illustration for Lemma 2.8 (picture not to scale).

See Figure 5, right.

Proof. Let $g = be \cap cf$, $h = ad \cap be$. Subdivide γ into $\gamma_1 = g_x - c_x$, $\gamma_2 = h_x - g_x$, $\gamma_3 = d_x - h_x$. By Lemma 2.5 we have

$$\frac{\alpha_1}{\beta_1} = \frac{\alpha_2 + \gamma_1 + \gamma_2}{\gamma_3}, \quad \frac{\alpha_2}{\beta_2} = \frac{\gamma_1}{\gamma_2 + \gamma_3 + \beta_1};$$

from which the first two claims follow.

By the first claim we have

$$\frac{\beta_1}{\beta_2} < \frac{\gamma}{\alpha_2} < \frac{\gamma + \beta_1 + \beta_2}{\alpha_2};$$

hence, if β_1/β_2 is larger than 1, then so is $(\gamma + \beta_1 + \beta_2)/\alpha_2$, implying the third claim. \square

Definition 2.7. Let s_1, s_2, \dots, s_m be segments whose endpoints appear in the order $L_{s_1} \dots L_{s_m} R_{s_1} \dots R_{s_m}$. These segments are called a *wide set* if their x -coordinates satisfy $R_{s_k x} - L_{s_1 x} > 2(R_{s_{k-1} x} - L_{s_1 x})$ for each $2 \leq k \leq m$.

Lemma 2.8. Let s_1, \dots, s_m be a wide set of segments that intersect concavely, and let $\alpha_k = L_{s_{k+1} x} - L_{s_k x}$ for $1 \leq k \leq m-1$. Then $\alpha_k > \alpha_{k+1} + \dots + \alpha_{m-1}$ for each $1 \leq k \leq m-2$.

Proof. Let $\gamma_k = R_{s_k x} - L_{s_1 x}$ for $1 \leq k \leq m$. See Figure 6. We are given that $\gamma_k > 2\gamma_{k-1}$ for each $2 \leq k \leq m$. Applying the first claim of Lemma 2.6 to segments s_k, s_{k+1}, s_m , we get

$$\frac{\alpha_k}{\alpha_{k+1} + \dots + \alpha_{m-1}} > \frac{\gamma_{k+1} - \gamma_k}{\gamma_k - \sum \alpha_i} > \frac{\gamma_{k+1} - \gamma_k}{\gamma_k} > 1.$$

The claim follows. \square




Figure 7: An impossible configuration of segments.

3 Warmup: A simple but useless impossible configuration

Theorem 3.1. *Let $a, b, c, d, e, 1, 2, 3, 4, 8, 9$ be eleven segments with endpoints on the parabola C , in left-to-right order*

$$L_8 L_1 L_a L_b L_2 R_8 L_c L_d R_1 R_2 L_e R_a L_3 L_4 R_b R_c L_9 R_3 R_d R_e R_4 R_9. \quad (1)$$

Then, it is impossible for segments 8, 1, 2 to intersect concavely, segments 3, 4, 9 to intersect concavely, and segments a, b, c, d, e to intersect concavely, all at the same time.

Proof. Suppose for a contradiction that $a, \dots, 9$ are segments satisfying all these properties. The intersection point of segments 1 and 2, which we shall call A , must lie left of R_8 , and the intersection point of segments 3 and 4, which we shall call B , must lie right of L_9 . See Figure 7.

Define:

$$\begin{aligned} \alpha_1 &= L_{bx} - L_{ax}, & \beta_1 &= R_{bx} - R_{ax}, \\ \alpha_2 &= L_{cx} - L_{bx}, & \beta_2 &= R_{cx} - R_{bx}, \\ \alpha_3 &= L_{dx} - L_{cx}, & \beta_3 &= R_{dx} - R_{cx}, \\ \alpha_4 &= L_{ex} - L_{dx}, & \beta_4 &= R_{ex} - R_{dx}. \end{aligned}$$

Segments a, b, c, d, e must intersect concavely, so by the second claim of Lemma 2.6, we must have

$$\frac{\alpha_1}{\beta_1} > \frac{\alpha_2}{\beta_2} > \frac{\alpha_3}{\beta_3} > \frac{\alpha_4}{\beta_4}. \quad (2)$$

We will show, however, that this is impossible. Define:

$$\begin{aligned} p &= L_{2x} - L_{1x}, & p' &= L_{4x} - L_{3x}, \\ r &= A_x - L_{2x}, & r' &= B_x - L_{4x}, \\ s &= R_{1x} - A_x, & s' &= R_{3x} - B_x, \\ q &= R_{2x} - R_{1x}, & q' &= R_{4x} - R_{3x}. \end{aligned}$$

Then,

$$\begin{aligned} \alpha_1 &< p, & \beta_1 &> p', \\ \alpha_2 &> r, & \beta_2 &< r', \\ \alpha_3 &< s, & \beta_3 &> s', \\ \alpha_4 &> q, & \beta_4 &< q'. \end{aligned}$$

Furthermore, by Lemma 2.5 we have $p/q = r/s$, $p'/q' = r'/s'$. Hence,

$$\frac{\alpha_1 \alpha_3}{\beta_1 \beta_3} < \frac{ps}{p's'} = \frac{qr}{q'r'} < \frac{\alpha_2 \alpha_4}{\beta_2 \beta_4},$$

contradicting (2). □

Corollary 3.2. *Let S be the lower inner-zone sequence of the parabola C in an arrangement of lines. Then S cannot contain a subsequence isomorphic to*

$$u = 81ab12181cd12dedcbab34bc49434de49.$$

Proof. We have

$$E(u) = L_8 L_1 L_a L_b L_2 R_8 L_c L_d R_1 R_2 L_e R_a L_3 L_4 R_b R_c L_9 R_3 R_d R_e R_4 R_9,$$

exactly matching (1). Furthermore, as a tedious examination shows, all symbols but 8 are left-clamped in u , and all symbols but 9 are right-clamped in u . Furthermore, u contains the N -shaped subsequences 8121812, 3494349, and $abcdedcbabcde$. Hence, by Lemma 2.2 and Observation 2.3, u forces the impossible segment configuration of Theorem 3.1. □

(We obtained the sequence u by simply taking the lower inner-zone sequence of the configuration of Theorem 3.1, and removing from it unnecessary symbols.)

If we are only interested in a pattern avoided by N , the lower-envelope sequence, then we can omit the symbols 8 and 9 from u . Their only role is preventing the intersection points A and B from “hiding” above the segment c .

Unfortunately, as we said in the Introduction, the forbidden pattern u is useless for establishing Conjecture 1.1, since u contains both $abcacbc$ and its reversal (e.g., $be4b44e4$, $1a11d1ad$). Furthermore, there does not seem to be a simple way to “fix” u .

4 A more promising impossible configuration

In this section we construct a 173-segment impossible configuration X . Then, in Section 5 we will show that the Hart–Sharir sequences eventually force the configuration X .

We will now work with endpoint sequences in which some contiguous subsequences (*blocks*) that contain only left endpoints are designated as *special blocks*. It will always be the case that all the special blocks in a sequence have the same length. We denote special blocks by enclosing them in parentheses.

We define an operation on endpoint sequences called *endpoint shuffling*. This operation is derived from the *shuffling* operation used to construct the Hart–Sharir sequences, which we will present in Section 5 below.

Let A be a sequence that has k special blocks of length m , and let B be a sequence that has ℓ special blocks of length k . Then the *endpoint shuffle* of A and B , denoted $A \circ B$, is a new sequence having $k\ell$ special blocks of length $m + 1$, formed as follows: We make ℓ copies of A (one for each special block of B), each one having “fresh” symbols that do not occur in B nor in any other copy of A .

For each special block $\Gamma_i = (L_1 \cdots L_k)$ in B , $1 \leq i \leq \ell$, let A_i be the i th copy of A . We insert each L_j at the end of the j th special block of A_i . Then we insert the resulting sequence in place of Γ_i in B . The result of all these replacements is the desired sequence $A \circ B$.

For example, let

$$A = (L_a) (L_b) (L_c) R_a R_b R_c, \quad B = (L_1 L_2 L_3) (L_4 L_5 L_6) R_1 R_4 R_2 R_5 R_3 R_6.$$

Then,

$$A \circ B = (L_a L_1) (L_b L_2) (L_c L_3) R_a R_b R_c (L_{a'} L_4) \\ (L_{b'} L_5) (L_{c'} L_6) R_{a'} R_{b'} R_{c'} R_1 R_4 R_2 R_5 R_3 R_6.$$

Now, define the following endpoint sequences:

$$F_m = (L_1 \cdots L_m) R_1 \cdots R_m, \quad m \geq 1; \\ Z_m = L_a L_b (L_1 \cdots L_m) R_1 \cdots R_m L_c R_a \\ (L_{m+1} \cdots L_{2m}) R_{m+1} \cdots R_{2m} R_b R_c, \quad m \geq 1; \\ Y = L_d L_e () () L_f R_d () () R_e R_f (),$$

where Y has five empty special blocks.

Recall from Section 2 the definition of a *wide* set of segments. We will now show how, by shuffling the Z and F sequences in the appropriate way, we can create, for every n , a configuration that contains a wide set of n segments.⁷

Consider the segment configuration

$$T_n = (\cdots (((Z_1 \circ Z_2) \circ Z_4) \circ Z_8) \circ \cdots Z_{2^{n-1}}) \circ F_{2^n}.$$

T_n contains 2^{n-k-1} copies of Z_{2^k} for each $0 \leq k < n$, plus one copy of F_{2^n} . Each copy of Z_m contains its own triple of segments a, b, c . These segments resemble a cat’s whiskers, so let

⁷With a slight abuse of terminology, we do not always distinguish between endpoint sequences and the corresponding segment configurations.




Figure 8: The configuration T_n contains *whiskers* (triples of segments labeled a, b, c) organized in a nested, binary-tree structure. In addition, it contains *numeric segments*, grouped into sets of size $n + 1$. This sketch shows only one such set of numeric segments, and it only shows their endpoints. For each set of numeric segments, their left endpoints lie in a special block of length $n + 1$ (indicated by parentheses in the figure).

us call each such triple a, b, c a *whisker* for short. The whiskers of T_n are nested inside one another in the form of a complete binary tree of height n .

In addition, each copy F_m and Z_m contains segments labeled $1, \dots, m$ and $m + 1, \dots, 2m$, which we shall call *numeric segments*. In total, T_n contains $(n + 1)2^n$ numeric segments, whose left endpoints appear in 2^n special blocks of length $n + 1$. Each such special block appears where we would expect to find the whiskers of level $n + 1$ of the binary tree.

Consider one such special block $(L^{(1)}, \dots, L^{(n+1)})$ (where the endpoints have been renamed for clarity of exposition). The corresponding right endpoints are located as follows: Let w_1, \dots, w_n be the whiskers containing our special block, where w_1 is the innermost whisker and w_n is the outermost one. Then, the first right endpoint $R^{(1)}$ appears within w_1 , right after $L^{(n+1)}$. For each $2 \leq i \leq n$, the right endpoint $R^{(i)}$ appears outside w_{i-1} but inside w_i . Finally, $R^{(n+1)}$ appears outside w_n . See Figure 8 for a rough sketch of T_n .

Lemma 4.1. *Suppose that, in T_n , the segments a, b, c in each whisker intersect concavely. Then at least one the sets of $n + 1$ numeric segments must be wide.*

Proof. Let $w = (a, b, c)$ be the outermost whisker of T_n . By the third claim of Lemma 2.6, we have $R_{bx} - R_{ax} < R_{cx} - R_{bx}$ or $L_{cx} - L_{bx} < R_{cx} - L_{cx}$ (or both). In the first case, we proceed to the left branch in the binary tree of whiskers. In the second case, we proceed to the right branch. Let $w' = (a', b', c')$ be the next whisker in the chosen branch. We invoke Lemma 2.6 again on w' , and we proceed in this way, going down the binary tree of whiskers until we reach a “leaf” whisker w^* . We invoke Lemma 2.6 one final time on w^* in the same way, and we choose one of the two special blocks that lie within w^* . The $n + 1$ numeric segments in that special block form a wide set, as can be easily checked. \square

Remark 4.2. As we saw, in the whisker $w = (a, b, c)$ of Z_m , one of two gaps is shorter than the distance from the end of the gap to the end of the whisker. If we could find a configuration in which a *specific* gap is shorter than the distance to the end of the configuration, we could obviate the need for a binary tree and significantly reduce the size of our construction.

Next, we show how a wide set of 5 segments can lead to trouble:

Lemma 4.3. *Consider the sequence*

$$Y \circ F_5 = L_d L_e L_1 L_2 L_f R_d L_3 L_4 R_e R_f L_5 R_1 R_2 R_3 R_4 R_5.$$

It is impossible for the segments d, e, f to intersect concavely, and for the segments $1, 2, 3, 4, 5$ to form a wide set and intersect concavely, all at the same time.

Proof. Suppose for a contradiction that we have a realization of such a configuration. Applying Lemma 2.8 on the segments $1, 2, 3, 4, 5$ (with $k = 1$ and $k = 3$), we have

$$L_{2x} - L_{1x} > L_{5x} - L_{2x}, \quad L_{4x} - L_{3x} > L_{5x} - L_{4x}.$$

Hence, we have both

$$R_{ex} - R_{dx} > L_{4x} - L_{3x} > L_{5x} - L_{4x} > R_{fx} - R_{ex}$$

and

$$L_{fx} - L_{ex} > L_{2x} - L_{1x} > L_{5x} - L_{2x} > R_{fx} - L_{fx},$$

contradicting the third claim of Lemma 2.6 on the segments d, e, f . \square

Now we can put everything together. Define the endpoint sequence

$$X = Y \circ T_4.$$

X contains 15 whiskers, 16 groups of segments of type d, e, f , and 16 groups of numeric segments, with five segments in each group. Hence, X contains a total of 173 segments.

Theorem 4.4. *It is impossible to realize X such that the segments a, b, c in each whisker intersect concavely, the segments d, e, f in each copy of Y intersect concavely, and the five numeric segments in each group of numeric segments intersect concavely.*

Proof. In X , each group of numeric segments of T_4 is shuffled into a copy of Y as in the premise of Lemma 4.3. But by Lemma 4.1, one of these groups must form a wide set. This is impossible by Lemma 4.3. \square

5 The Hart–Sharir sequences are unrealizable

In this section we show that the Hart–Sharir sequences eventually force a copy of the configuration X whose segments intersect in the way specified in Theorem 4.4. (Recall the definition of *forcing* a segment configuration from Section 1.6.) Hence, the Hart–Sharir sequences cannot be realized as lower inner-zone sequences of a parabola.

5.1 The Hart–Sharir sequences

We first recall the Hart–Sharir construction [9] of superlinear-length order-3 DS sequences.

The Hart–Sharir sequences form a two-dimensional array $S_k(m)$, for $k, m \geq 1$; they satisfy the following properties:

- Some contiguous subsequences (*blocks*) in $S_k(m)$ are designated as *special blocks*.
- All special blocks in $S_k(m)$ have length exactly m .
- Each symbol in $S_k(m)$ makes its first occurrence in a special block, and each special block contains only first occurrences of symbols.
- $S_k(m)$ contains no adjacent repetitions and no alternation *ababa*.
- Each symbol in $S_k(m)$ occurs at least twice (unless $k = m = 1$).

The construction uses an operation called *shuffling*, which, as we will see, is very closely related to the *endpoint shuffling* defined in Section 4.

Definition 5.1. Let A be a sequence that has k special blocks of length m , and let B be a sequence that has ℓ special blocks of length k . Then the *shuffle* of A and B , denoted $A \bullet B$, is a new sequence having $k\ell$ special blocks of length $m + 1$, formed as follows: We make ℓ copies of A (one for each special block of B), each one having “fresh” symbols that do not occur in B or in any other copy of A .

For each special block $\Gamma_i = (a_1 a_2 \cdots a_k)$ in B , $1 \leq i \leq \ell$, let A_i be the i th copy of A . For each special block Δ_j in A_i , $1 \leq j \leq k$, we insert the symbol a_j at the end of Δ_j (so its length grows from m to $m + 1$) and we duplicate the m th symbol of Δ_j immediately after Δ_j . Then we place another copy of a_k immediately after A_i . Call the resulting sequence A'_i .

Then $A \bullet B$ is obtained from B by replacing each special block Γ_i in it by A'_i .

In the construction of $A \bullet B$, the symbols of the copies of A are usually called *local*, while the symbols of B are called *global*.

For example, let $A = (a)(b)(c)babcb$ and $B = (123)21(456)5414525636$. Then,

$$A \bullet B = (a1)a(b2)b(c3)cbabc3 \ 21 \ (a'4)a'(b'5)b'(c'6)c'b'a'b'c'6 \ 5414525636.$$

Let S be a sequence of symbols in which the first occurrences of the symbols appear in special blocks, and in which each symbol occurs at least twice. Let us extend the definition of the *endpoint sequence* $E(S)$ from Section 2, by specifying that $E(S)$ has special blocks of left endpoints, which are inherited from the special blocks of S in the natural way.

Observation 5.2. *The relation between shuffling and endpoint shuffling is as follows:*

$$E(A \bullet B) = E(A) \circ E(B).$$

Proof. This is immediate from the definitions of \bullet and \circ . \square

The Hart–Sharir sequences $S_k(m)$ are defined by double induction. The base cases are $k = 1$ and $m = 1$. For $k = 1$ we let

$$S_1(m) = (12 \cdots (m-1)m) (m-1) \cdots 212 \cdots m,$$

be an N -shaped sequence having a single special block of length m . (Actually, of all sequences $S_1(m)$, the construction only uses $S_1(2) = (12)12$.)

For $m = 1$, $k \geq 2$, we let $S_k(1)$ be equal to $S_{k-1}(2)$, but with each special block of size 2 split into two special blocks of size 1.

Finally, for $m, k \geq 2$, we let

$$S_k(m) = S_k(m-1) \bullet S_{k-1}(N),$$

where N is the number of special blocks in $S_k(m-1)$.

Thus, we have, up to a renaming of the symbols,

$$\begin{aligned} S_2(1) &= (1)(2)12, \\ S_2(2) &= (12)1(34)313424, \\ S_2(3) &= (123)21(456)5414525636, \\ S_2(4) &= (1234)321(5678)76515626737848, \\ &\vdots \\ S_3(1) &= (1)(2)1(3)(4)313424, \\ S_3(2) &= (12)1(34)31(56)5(78)75157378642 \\ &\quad (9A)9(BC)B9(DE)D(FG)FD9DFBFGEC A2AC4CE6EG8G, \\ &\vdots \\ S_4(1) &= (1)(2)1(3)(4)31(5)(6)5(7)(8)75157378642 \\ &\quad (9)(A)9(B)(C)B9(D)(E)D(F)(G)FD9DFBFGEC A2AC4CE6EG8G, \\ &\vdots \end{aligned}$$

Note that, in the construction of $S_k(m)$, the special blocks of $S_{k-1}(N)$ “dissolve”, and the only special blocks present in $S_k(m)$ are those that come from the copies of $S_k(m-1)$ (enlarged by one).

5.2 Properties

Here we establish some important properties of $S_k(m)$.

Lemma 5.3. *Each special block $(1 \cdots m)$ in $S_k(m)$ is immediately followed by $(m-1) \cdots 1$, and followed later on by $\cdots 2 \cdots 3 \cdots \cdots m$, thus forming an N -shaped subsequence.*

Proof. By induction. The claim is true for $k = 1$ and all m . If the claim is true for $S_{k-1}(2)$, then it is also true for $S_k(1)$. Now, let $k, m \geq 2$, and suppose the claim is true for $S_k(m-1)$. Then, by the definition of \bullet , and since each symbol in $S_{k-1}(N)$ occurs at least twice, the claim is also true for $S_k(m) = S_k(m-1) \bullet S_{k-1}(N)$. \square

Lemma 5.4. $S_k(m)$ does not contain $ababa$ (Hart, Sharir [9]) nor $abcacbc$ (Pettie [19]).⁸

Proof sketch. For the first claim, suppose for a contradiction that k and m are minimal such that $S_k(m)$ contains $ababa$. Recall that $S_k(m) = S_k(m-1) \bullet S_{k-1}(N)$. Each of the symbols a, b must have come either from a copy of $S_k(m-1)$ (in which case it is a local symbol) or from $S_{k-1}(N)$ (in which case it is a global symbol). A case analysis shows that none of the possibilities work. For example, it cannot be that a is local and b is global, because then there would be at most one b between two a 's. It cannot be either that a is global and b is local, because then only the *first* a could appear between two b 's.

For the second claim, first note that Lemma 5.3 implies that $S_k(m)$ cannot contain $(bc)ccb$ (with the first b and c in the same special block), since, otherwise, $S_k(m)$ would contain $bcbcb$. For the same reason, if $S_k(m)$ contains $(bc)cb$, then c is *not* the last symbol in the special block.

Now, suppose for a contradiction that k, m are minimal such that $S_k(m)$ contains $abcacbc$. There are eight possibilities as to which of the symbols a, b, c are local and which are global. We rule them all out by a case analysis. First of all, if a, b, c are all local (resp. all global), then the only possibility would be for $S_k(m-1)$ (resp. $S_{k-1}(N)$) to contain $abcacbc$, so that the second c gets duplicated into cc in $S_k(m)$. But this cannot happen, since only *first* occurrences of symbols get duplicated. The case of a being local and b, c being global is ruled out by the considerations in the previous paragraph. The remaining cases are also readily ruled out. \square

Definition 5.5. Let A and B be two sequences in which the first occurrences of symbols appear in special blocks. Then we say that B *structurally contains* A if B contains a subsequence A' that not only is isomorphic to A , but for every two symbols in A' , their first occurrences lie in the same special block if and only if the corresponding symbols in A lie in the same special block.

We now want to prove that $S_{k'}(m')$ structurally contains $S_k(m)$ for all $k' \geq k$, $m' \geq m$. Hence, any problematic configuration that arises in some $S_k(m)$ will also be present in all subsequent $S_{k'}(m')$.

The proof is somewhat delicate. It is clear, for example, that $S_{k+1}(m)$ contains a sequence isomorphic to $S_k(m)$: $S_{k+1}(m)$ contains a global copy of $S_k(N)$ for some N larger than m , and in turn $S_k(N)$ contains many local copies of $S_k(N-1)$, et cetera. This simple observation, however, does not imply that $S_{k+1}(m)$ *structurally* contains $S_k(m)$, because the copy of $S_k(m)$ that we found in $S_{k+1}(m)$ has its special blocks completely “dissolved”. To prove that $S_{k+1}(m)$ structurally contains $S_k(m)$ we have to work a bit more carefully.

Definition 5.6. The *rank* of a symbol a in $S_k(m)$ is the position (between 1 and m) that the first occurrence of a occupies within its special block.

Thus, the local symbols of $S_k(m)$ are those with ranks $1, \dots, m-1$, and the global symbols are those with rank m .

Lemma 5.7. For every $k' \geq k$, $m' \geq m$, and for every choice of m ranks $1 \leq r_1 < r_2 < \dots < r_m \leq m'$, the sequence $S_{k'}(m')$ structurally contains $S_k(m)$ using symbols of these ranks.

⁸Note that $S_3(2)$ already contains $(abcacbc)^R$.

Proof. Denote $B = S_k(m)$ and $D = S_{k'}(m')$, and let N_2 and N'_2 denote, respectively, the number of special blocks in B and D . Note that, in order to specify how B lies within D , all we have to do is specify which N_2 special blocks $1 \leq b_1 < b_2 < \dots < b_{N_2} \leq N'_2$ of D we take the symbols from.

We will first take care of the base cases $k = 1$ and $m = 1$.

If $k = 1$ then the claim follows by Lemma 5.3. If $m = m' = 1$ then there is nothing to prove, since in this case structural containment is the same as regular containment. Now suppose $m = 1$ and $m' \geq 2$. Recall that in this case we have $B = S_{k-1}(2)$. We are given a rank r_1 . The symbols of rank r_1 in D are the global symbols of $S_{k'}(r_1)$. The global sequence used in forming $S_{k'}(r_1)$ is $S_{k'-1}(N)$ for some N . This latter sequence contains $B = S_{k-1}(2)$ in the regular sense, which is enough for us.

Now suppose $k, m \geq 2$.

For convenience let $A = S_k(m-1)$, $C = S_{k'}(m'-1)$. Let N and N' be, respectively, the number of special blocks in A and C . Let $X = S_{k-1}(N)$, $Y = S_{k'-1}(N')$. Hence,

$$\begin{aligned} B &= A \bullet X, \\ D &= C \bullet Y. \end{aligned}$$

We assume by induction that C contains A , and that Y contains X , in the stronger sense of our lemma. If $m' > m$ we can also assume by induction that C contains B in this stronger sense. Our objective, as we said, is to show that D contains B in this stronger sense.

We consider two cases: If $r_m < m'$ (which implies $m < m'$), then we wish to find B using only local symbols of D . But we know by induction that C (which is structurally contained in D) contains B in the desired way.

The second case is if $r_m = m'$. We know by induction that C contains A using ranks r_1, \dots, r_{m-1} . Let $1 \leq b_1 < b_2 < \dots < b_N \leq N'$ be the special blocks of C from which we take the symbols that form A this way.

Next, find X in Y using ranks b_1, b_2, \dots, b_N . We know this is possible by induction. Let $1 \leq c_1 < c_2 < \dots$ be the special blocks of Y from which we take the symbols that form X this way.

Now, in order to find B in D , take only the local copies numbered c_1, c_2, \dots of C , and within each local copy of C , take only the special blocks numbered $b_1, b_2, \dots, b_{N'}$. Then the symbols of Y that are shuffled into these special blocks are exactly those that form X .

Hence, we exactly mimic the construction of B from A and X inside the construction of D from C and Y —with only one exception: In the construction of B , we duplicate the last symbols of the special blocks of A and X . This might not happen to the corresponding symbols in D , if these symbols are not the last symbols of the special blocks of C and Y . However, we do not need these duplications, since they are *already* present immediately after the special blocks, by Lemma 5.3. \square

5.3 Clamped symbols in the Hart–Sharir sequences

Recall the definition of left- and right-clamped symbols (Definition 2.1) and its use in Lemma 2.2. We now proceed to show that most symbols in $S_k(m)$ are left- and right-clamped.

Lemma 5.8. *All symbols in $S_k(m)$ that have rank at least 2 are left-clamped, and all the symbols, other than the very last symbol in $S_k(m)$, are right-clamped.*

Proof. The first claim follows immediately from Lemma 5.3.

For the second claim, recall that $S_k(m) = S_k(m-1) \bullet S_{k-1}(N)$. If a is the very last symbol of $S_{k-1}(N)$, then it is the very last symbol of $S_k(m)$ and there is nothing to prove. If a is the very last symbol of a copy S' of $S_k(m-1)$, then its last occurrence in $S_k(m)$ is immediately followed by a global symbol b , whose first occurrence was shuffled at the end of the *last* special block of S' . The first occurrence of a must lie further to the left (also in S').

In any other case, the claim follows by induction on $S_k(m-1)$ or $S_{k-1}(N)$, depending on whether a is a local or a global symbol, since the shuffling operation does not separate the last occurrence of any symbol from the immediately following symbol. \square

Corollary 5.9. *Once a copy of $S_k(m)$ is shuffled into another sequence $(S_{k+1}(m'))$ for some $m' > 1$, all its symbols are left-clamped, and all its symbols except for the very last one are right-clamped.*

Proof. All the symbols of the copy of $S_k(m)$ that resides in $S_{k+1}(m')$ are global (their rank is m'). \square

We can also say something about left-clamped rank-1 symbols:

Lemma 5.10. *Let $k, m \geq 2$, and let a be a rank-1 symbol in $S_k(m) = S_k(m-1) \bullet S_{k-1}(N)$. If a was left-clamped in its local copy of $S_k(m-1)$, then it is also left-clamped in $S_k(m)$.*

Proof. By Lemma 5.3, $S_k(m-1)$ contains adjacent special blocks only if $m = 2$. Hence, for $m \geq 3$, the symbol b immediately preceding the first a in $S_k(m-1)$ does not belong to a special block, so b remains adjacent to that a in $S_k(m)$. For $m = 2$, the said symbol b might belong to a special block. But then, in the construction of $S_k(m)$, a copy of b is created and placed right before the first a . \square

5.4 Geometric unrealizability

We are now ready to show that the Hart–Sharir sequences force the impossible configuration X of Section 4.

Theorem 5.11. *There exists an m such that the Hart–Sharir sequence $S_7(m)$ forces a copy of the configuration X whose segments intersect in the way specified in Theorem 4.4. Hence, if S is the lower inner-zone sequence of the parabola C , then S cannot contain $S_7(m)$ as a subsequence.⁹*

Proof. Recall that in Section 4 we defined

$$F_m = (L_1 \cdots L_m) \quad R_1 \cdots R_m.$$

Let

$$Z_{j,m} = L_{a_1} F_m L_{a_2} F_m \cdots F_m L_{a_j} \quad R_{a_1} F_m R_{a_2} F_m \cdots F_m R_{a_j}$$

(where the $2j-2$ copies of F_m use distinct symbols). Note that the Z_m of Section 4 is contained in $Z_{j,m}$ for $j \geq 3$. Following this correspondence, let us call the j -tuple of segments a_1, \dots, a_j in $Z_{j,m}$ a *whisker*, and let us call the segments in the copies of F_m *numeric segments*.

⁹It should probably be enough to take $m = 2$ in the theorem, but the calculations do not seem to be worth the effort.

Define

$$T_{j,n} = (\cdots (((Z_{j,1} \circ Z_{j,2j-2}) \circ Z_{j,(2j-2)^2}) \circ Z_{j,(2j-2)^3}) \circ \cdots Z_{j,(2j-2)^{n-1}}) \circ F_{(2j-2)^n}.$$

For $j \geq 3$, $T_{j,n}$ is a supersequence of the endpoint sequence T_n that we defined in Section 4. $T_{j,n}$ contains whiskers that are nested inside one another in the form of a full $(2j-2)$ -ary tree of height n . In addition, it contains numeric segments, whose left endpoints are grouped into special blocks of length $n+1$, and whose right endpoints are located in a manner analogous to the one described in Section 4.

Our first goal is to show that the Hart–Sharir sequences force segment configurations of the form $T_{j,n}$ for arbitrarily large j and n , in which the j segments in each whisker intersect concavely. From Lemma 5.3 and Observation 2.3, it follows that the $n+1$ numeric segments in each special block must also intersect concavely.

Observation 5.12. $E(S_1(m)) = F_m$, and $E(S_2(m))$ contains

$$(L_1 \cdots L_m) () R_1 \cdots R_m. \quad (3)$$

Corollary 5.13. $E(S_3(m))$ contains

$$(L_1 L_2 \cdots L_m) () R_1 () R_2 () \cdots () R_m. \quad (4)$$

Proof. By induction on m . Consider $S_3(m) = S_3(m-1) \bullet S_2(N)$, and assume by induction that $S_3(m-1)$ contains an instance A of (4) with $m-1$ in place of m . Consider an instance B of (3) in $E(S_2(N))$ with N in place of m . The first special block of B is shuffled into a copy of $S_3(m-1)$; hence, one of its symbols L^* is inserted at the end of the first special block of A in this copy of $S_3(m-1)$. Furthermore, the $()$ and the corresponding symbol R^* of B are placed after this copy of A . Hence, $S_3(m)$ contains (4). \square

Corollary 5.14. For each $m \geq 2$, $S_4(m)$ forces a configuration of the form

$$L_1 () L_2 () \cdots () L_N () R_1 () R_2 () \cdots () R_N \quad (5)$$

for some very large $N = N(m)$, in which the segments $1, 2, \dots, N$ intersect concavely and have rank m .

Proof. The global sequence $S_3(N')$ used in forming $S_4(m)$ satisfies Corollary 5.13. The left endpoints L_1, L_2, \dots of (4) receive rank m and go into separate special blocks. In order to guarantee the presence of a special block between L_i and L_{i+1} for each i , we “sacrifice” every second symbol among $1, 2, \dots, N'$; we are still left with $N = N'/2$ symbols. The N -shaped sequence $1 \cdots N' \cdots 1 \cdots N'$ that was present in $S_3(N')$ (by Lemma 5.3) is obviously still present in $S_4(m)$. \square

Corollary 5.15. For each $m \geq 2$, $S_5(m)$ forces a configuration of the form $Z_{N,m-1}$ for some very large $N = N(m)$, in which the whisker segments a_1, \dots, a_N intersect concavely, and in which the numeric segments of each copy of F_{m-1} have ranks $1, \dots, m-1$.

Proof. The global sequence $S_4(N')$ used in forming $S_5(m)$ satisfies Corollary 5.14. Each special block of (5) is replaced by a copy of $S_5(m-1)$, which structurally contains $S_1(m-1)$. \square

Finally:

Corollary 5.16. *For every j and n , if m is large enough, then $S_6(m)$ forces a configuration of the form $T_{j,n}$, in which the segments a_1, \dots, a_j in each whisker intersect concavely.*

Proof. The iterated shuffling used to form $T_{j,n}$ occurs naturally in the formation of $S_6(m) = S_6(m-1) \bullet S_5(N)$.

More precisely, we will show by induction on n that, if m is large enough in terms of j and n , then $S_6(m)$ forces a configuration of the form

$$T'_{j,n} = (\cdots ((Z_{j,1} \circ Z_{j,2j-2}) \circ Z_{j,(2j-2)^2}) \circ \cdots) \circ Z_{j,(2j-2)^n},$$

in which the segments of each whisker intersect concavely. Since $Z_{j,(2j-2)^n}$ contains $F_{(2j-2)^n}$, it will follow that $T'_{j,n}$ contains $T_{j,n}$.

Suppose by induction that, for some m , the sequence $S_6(m)$ forces such a configuration $T'_{j,n-1}$. Let N be the number of special blocks of $S_6(m)$, and let $1 \leq b_1 < b_2 < \cdots < b_k \leq N$, for $k = (2j-2)^n$, be the special blocks of $S_6(m)$ involved in this occurrence of $T'_{j,n-1}$. By Corollary 5.15, there exists an N' such that $S_5(N')$ forces a copy of Z_{j,b_k} in which the numeric segments of each F_{b_k} have ranks $1, \dots, b_k$. Out of the b_k segments in each F_{b_k} , we are only interested in the ones ranked b_1, \dots, b_k : They are the ones that will be shuffled into the right places in $T'_{j,n-1}$. Let $m' \geq m$ and $N'' \geq N'$ be such that $S_6(m'+1) = S_6(m') \bullet S_5(N'')$. Now, $S_6(m')$ contains a copy of $S_6(m)$ using the first N special blocks of $S_6(m')$. And the above-mentioned copy of Z_{j,b_k} that is present in $S_5(N')$ is also present in $S_5(N'')$. Hence, $S_6(m'+1)$ forces the desired copy of $T'_{j,n}$. \square

Hence, as claimed, the sequences $S_6(m)$ force arbitrarily wide and tall configurations of the form $T_{j,n}$, in which the segments of each whisker intersect concavely. In particular, they force one such copy of $T_{3,4}$, which contains the configuration T_4 of Section 4.

Our second goal is to show that this sequence $T_{3,4}$ is appropriately shuffled into a sequence containing Y (which was defined in Section 4).

We observe that (5) already contains Y whenever $N \geq 6$. Hence, $S_7(1)$ also contains Y . Let $b_1 < b_2 < \cdots < b_5$ be the special blocks of $S_7(1)$ involved in this occurrence of Y . By Corollary 5.16 and Lemma 5.7, $S_6(m)$ contains, for large enough m , a copy of $T_{3,4}$ in which the numeric segments have ranks b_1, b_2, \dots, b_5 . Therefore, there exists an m (probably $m = 2$ should be enough) such that, in $S_7(m) = S_7(m-1) \bullet S_6(N)$, these numeric segments are shuffled into the right places, creating the desired copy of X . In this copy of X , all the sets of segments that should intersect concavely according to Theorem 4.4, actually do.

Finally, all the symbols in this copy of X are left- and right-clamped in $S_7(m)$: The copy of Y in $S_7(1) = S_6(2)$ uses symbols that, in $S_6(2)$, had rank 2, so they were clamped by Lemma 5.8. Hence, by Lemma 5.10, they stay clamped in $S_7(m)$. And the symbols of $S_6(N)$ become global in $S_7(2)$, so they are also clamped by Lemma 5.8. Hence, by Lemma 2.2, if S is a lower inner-zone sequence that contains $S_7(m)$, then the endpoints of X appear in the right order in $E(S)$. \square

5.5 How we found these results

Our results of Sections 4 and 5 were obtained as follows:

Matoušek [15] and Sharir and Agarwal [24] describe a construction by P. Shor of segments in the plane whose lower-envelope sequences are the Hart–Sharir sequences. We tried to force the segment endpoints in the construction to lie on the parabola C . We managed to do this for $S_1(m)$, $S_2(m)$, and $S_3(m)$ and all m , but for $S_4(m)$ this seems impossible.

Shor’s construction is based on *fans*—sets of segments that intersect concavely, whose left endpoints are very close to one another, and whose lengths increase very rapidly. Global fans have their left endpoints shuffled into tiny local fans. In order for the global fan to intersect concavely, its segments are given slopes $1, 1 + \varepsilon_1, 1 + \varepsilon_2, \dots$ for very small values of $\varepsilon_1, \varepsilon_2, \dots$. This gives a lot of freedom to play with the exact position of the segments’ left endpoints.

However, if we want all endpoints to lie on a parabola, then the slopes in the global fan must increase very rapidly, which leads to the absurd requirement that the distances between the left endpoints decrease very rapidly (Lemma 2.8 above). Then it is impossible to appropriately shuffle the global fan into the local fans.

The impossible segment configuration X of Section 4 (which, as we saw in Section 5, is forced by $S_7(m)$) is the best way we found to isolate the contradiction.

It would be nice to be able to isolate a smaller impossible segment configuration forced by the Hart–Sharir sequences (say, by $S_4(m)$). However, it is unlikely that such an improvement would be of additional help in proving Conjecture 1.1. The most promising line of attack is described in the next section.

6 Directions for future work

We believe that our geometric results are sufficient to prove Conjecture 1.1, and that the remaining work is purely combinatorial:

Conjecture 6.1. *The Hart–Sharir sequences are the only way to achieve superlinear-length $ababa$ -free sequences. Namely, for every Hart–Sharir sequence $S_k(m)$ we have*

$$\text{Ex}(\{ababa, S_k(m), (S_k(m))^R\}, n) = O(n);$$

where the hidden constant depends on k and m .

In order to establish Conjecture 1.1, it is enough to prove Conjecture 6.1 for the specific (gigantic) case of Theorem 5.11. However, our hope is that Conjecture 6.1 can be somehow more easily proven for *all* k and m by a double induction argument.

Conjecture 6.1 is known to be true for $k = 1$, since $S_1(m)$ are N -shaped sequences: Klazar and Valtr [14] showed that (even without forbidding $ababa$) we have $\text{Ex}(S_1(m), n) \leq c_m n$ for some constants c_m . Pettie [18] subsequently improved the dependence of c_m on m to $c_m \leq 2^{\Theta(m^2)}$, which is still quite large. No interesting lower bounds for c_m are known. In any case, we conjecture that, forbidding both $ababa$ and an N -shaped pattern, we should have $\text{Ex}(\{ababa, S_1(m)\}, n) \leq c'_m n$ for some quite small c'_m .

The first open case in Conjecture 6.1 is $k = m = 2$. In this case, $S_2(2) = abacdcacdbd \equiv (S_2(2))^R$; see Figure 9. However, as we mentioned in the Introduction, even the weaker conjecture, that $\text{Ex}(\{ababa, abcacbc\}, n) = O(n)$, is still open.

(Similarly, the other conjecture mentioned in the Introduction, namely that $\text{Ex}(\{ababa, abcacbc, (abcacbc)^R\}, n) = O(n)$, would follow from the case $k = 3, m = 2$ of Conjecture 6.1.)




Figure 9: The case $k = m = 2$ of Conjecture 6.1 states that this pattern must be present in any configuration of n x -monotone pseudosegments, if its lower-envelope sequence has length cn for some large enough constant c . The highlighted subsegments must be visible from $-\infty$.

6.1 Linear vs. nonlinear forbidden patterns

Conjecture 6.1 fits into the following more general question: For which patterns u is $\text{Ex}(u, n)$ linear in n ? This question has been previously explored in several papers [1, 11, 13, 18, 19]. The known results in this area are somewhat patchy, and a proof (or disproof) of our conjecture will shed additional light in this area. For one, our conjecture already highlights the fact that forbidding a *set* of patterns might have a stronger effect than forbidding each pattern separately, and hence, the right question should be: For which *sets* of patterns U is $\text{Ex}(U, n)$ linear in n ?

7 Conclusion

We end by listing some related open problems:

- As mentioned at the end of Section 5, it would be nice to isolate a smaller impossible segment configuration forced by the Hart–Sharir sequences (say, by $S_4(m)$ for some m); see Remark 4.2 above. (Still, it is unlikely that such an improvement would make proving Conjecture 1.1 much easier.)
- What if we do not require C to be a circle, but only a convex curve? It still seems impossible to implement Shor’s construction forcing the endpoints to lie on a convex curve.
- Can the *ababa*-free sequences of Nivasch–Geneson [8, 16] be realized as lower-envelope sequences of line segments? Is there an *ababa*-free sequence that *cannot* be realized in such a way?
- The longest Davenport–Schinzel sequences of order 4 (*ababab*-free) have length $\Theta(n \cdot 2^{\alpha(n)})$. However, no one knows how to realize them as lower-envelope sequences of parabolic segments. Perhaps it is impossible. One could start by finding forbidden patterns here.
- *Higher dimensions*: Raz [23] recently proved that the combinatorial complexity of the outer zone of the boundary of a convex body in an arrangement of hyperplanes in R^d is $O(n^{d-1})$. The complexity of the inner zone is only known to be $O(n^{d-1} \log n)$ (Aronov et al. [2]). Whether the latter is also linear in n is an open question.

Acknowledgements I would like to give special thanks to Seth Pettie for useful discussions on generalized DS sequences with different forbidden patterns. Thanks also to the referees of previous versions of this work for their careful reading and for their useful suggestions (including the idea in Section 1.4). Finally, thanks to Micha Sharir for useful discussions, and to Dan Halperin for encouraging me to work on this problem (several years ago).

References

- [1] R. Adamec, M. Klazar, and P. Valtr, Generalized Davenport–Schinzel sequences with linear upper bound, *Discrete Math.*, 108:219–229, 1992.
- [2] B. Aronov, M. Pellegrini, and M. Sharir, On the zone of a surface in a hyperplane arrangement, *Discrete Comput. Geom.*, 9:177–186, 1993.
- [3] M. Bern, D. Eppstein, P. Plassmann, and F. Yao, Horizon theorems for lines and polygons, in (Goodman et al., eds.) *Discrete and computational geometry: Papers from the DIMACS special year*, Vol. 6 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, AMS, 1991.
- [4] B. Chazelle, L. Guibas, and D. T. Lee, The power of geometric duality, *BIT* 25:76–90, 1985.
- [5] J. Cibulka and J. Kynčl, Tight bounds on the maximum size of a set of permutations with bounded VC-dimension, *J. Comb. Theory A*, 119:1461–1478, 2012.
- [6] H. Edelsbrunner, L. Guibas, J. Pach, R. Pollack, R. Seidel, and M. Sharir, Arrangements of curves in the plane—topology, combinatorics, and algorithms, *Theor. Comput. Sci.*, 92:319–336, 1992.
- [7] J. Fox, J. Pach, and A. Suk, The number of edges in k -quasi-planar graphs, *SIAM J. Discrete Math.*, 27:550–561, 2013.
- [8] J. Geneson, A relationship between generalized Davenport–Schinzel sequences and interval chains, *Electron. J. Comb.*, 22, P3.19, 10 pp., 2015.
- [9] S. Hart and M. Sharir, Nonlinearity of Davenport–Schinzel sequences and of generalized path compression schemes, *Combinatorica*, 6:151–177, 1986.
- [10] M. Klazar, A general upper bound in extremal theory of sequences, *Comment. Math. Univ. Carol.*, 33:737–746, 1992.
- [11] M. Klazar, Two results on a partial ordering of finite sequences, *Comment. Math. Univ. Carol.*, 34:697–705, 1993.
- [12] M. Klazar, On the maximum lengths of Davenport–Schinzel sequences, in (Graham et al., eds.) *Contemporary Trends in Discrete Mathematics*, Vol. 49 of *DIMACS Series in Discrete Mathematics and Theoretical Computer Science*, AMS, 1999.
- [13] M. Klazar, Generalized Davenport–Schinzel sequences: Results, problems, and applications, *Integers*, 2, A11, 39 pp., 2002.

- [14] M. Klazar and P. Valtr, Generalized Davenport–Schinzel sequences, *Combinatorica*, 14:463–476, 1994.
- [15] J. Matoušek, *Lectures on Discrete Geometry*, Springer, 2002.
- [16] G. Nivasch, Improved bounds and new techniques for Davenport–Schinzel sequences and their generalizations, *J. ACM*, 57, article 17, 44 pp., 2010.
- [17] S. Pettie, Splay trees, Davenport–Schinzel sequences, and the deque conjecture, *Proc. 19th Annu. ACM-SIAM Symp. Discrete Alg. (SODA 2008)*, pp. 1115–1124.
- [18] S. Pettie, On the structure and composition of forbidden sequences, with geometric applications, *Proc. 27th Annu. Symp. Comput. Geom (SoCG 2011)*, pp. 370–379.
- [19] S. Pettie, Origins of nonlinearity in Davenport–Schinzel sequences, *SIAM J. Discrete Math.*, 25:211–233, 2011.
- [20] S. Pettie, Generalized Davenport–Schinzel sequences and their 0-1 matrix counterparts, *J. Comb. Theory A*, 118:1863–1895, 2011.
- [21] S. Pettie, Sharp bounds on Davenport–Schinzel sequences of every order, *J. ACM*, 62, article 36, 40 pp., 2015.
- [22] S. Pettie, Three generalizations of Davenport–Schinzel sequences, *SIAM J. Discrete Math.*, 29:2189–2238, 2015.
- [23] O. E. Raz, On the zone of the boundary of a convex body, *Comp. Geom. Theor. Appl.*, 48:333–341, 2015.
- [24] M. Sharir and P. K. Agarwal, *Davenport–Schinzel Sequences and their Geometric Applications*, Cambridge University Press, 1995.
- [25] A. Suk and B. Walczak, New bounds on the maximum number of edges in k -quasi-planar graphs, *Comp. Geom. Theor. Appl.*, 50:24–33, 2015.
- [26] P. Valtr, Graph drawing with no k pairwise crossing edges, *5th Int. Symp. Graph Drawing*, 1997; *LNCS* 1353:205–218, 1997.
- [27] A. Wiernik and M. Sharir, Planar realizations of nonlinear Davenport–Schinzel sequences by segments, *Discrete Comput. Geom.*, 3:15–47, 1988.
- [28] I. M. Yaglom, *A Simple Non-Euclidean Geometry and its Physical Basis: An Elementary Account of Galilean Geometry and the Galilean Principle of Relativity*, Springer-Verlag, 1979.